arXiv:1610.07796v2 [cs.CL] 26 Oct 2016

Still not there? Comparing Traditional Sequence-to-Sequence Models to Encoder-Decoder Neural Networks on Monotone String Translation Tasks

Carsten Schnober[†‡], **Steffen Eger**[†], **Erik-Lân Do Dinh**[†], **and Iryna Gurevych**[†‡]

[†]Ubiquitous Knowledge Processing Lab (UKP-TUDA)
Department of Computer Science, Technische Universität Darmstadt
[‡]Ubiquitous Knowledge Processing Lab (UKP-DIPF)
German Institute for Educational Research and Educational Information
https://www.ukp.tu-darmstadt.de/

Abstract

We analyze the performance of encoder-decoder neural models and compare them with well-known established methods. The latter represent different classes of traditional approaches that are applied to the monotone sequence-to-sequence tasks OCR post-correction, spelling correction, grapheme-to-phoneme conversion, and lemmatization. Such tasks are of practical relevance for various higher-level research fields including *digital humanities*, automatic text correction, and speech recognition. We investigate how well generic deep-learning approaches adapt to these tasks, and how they perform in comparison with established and more specialized methods, including our own adaptation of pruned CRFs.

1 Introduction

Encoder-decoder neural models (Sutskever et al., 2014) are a generic deep-learning approach to sequence-to-sequence translation (Seq2Seq) tasks. They encode an input sequence into a vector representation from which the decoder generates an output. These models have shown to achieve state-of-the-art or at least highly competitive results for various NLP tasks including machine translation (Cho et al., 2014), conversation modeling (Vinyals and Le, 2015), question answering (Yin et al., 2016), and, more generally, language correction (Schmaltz et al., 2016; Xie et al., 2016).

We have noticed that, given the enormous interest currently surrounding neural architectures, recent research appears to somewhat over-enthusiastically praise the performance of encoder-decoder approaches for Seq2Seq tasks. For example, while the encoder-decoder G2P model by Rao et al. (2015) achieves an extremely low error rate on the CMUdict dataset (Kominek and Black, 2004), the neural architecture itself has a mediocre performance and only outperforms traditional models *in combination* with a weighted finite state transducer. Similarly, Faruqui et al. (2016) report on "par or better" performance of their inflection generation neural architecture. However, a closer inspection of their results suggests that their system is sometimes worse and sometimes better than traditional approaches.

Here, we aim for a more balanced comparison on three exemplary monotone[1] Seq2Seq tasks, namely *spelling correction*, *G2P conversion*, and *lemmatization*. Monotone Seq2Seq tasks such as morphological analysis/lemmatization, grapheme-to-phoneme conversion (G2P) (Yao and Zweig, 2015; Rao et al., 2015), transliteration (Sherif and Kondrak, 2007), and spelling correction (Brill and Moore, 2000) have been fundamental problem classes in natural language processing (NLP) ever since the origins of the field. Their simplicity vis-à-vis non-monotonic problems such as machine translation renders them as particularly tractable testbeds of technological progress. Unlike previous work, which has typically focussed on only one specific subproblem of monotone Seq2Seq tasks at a time, we consider model performances on *three* such tasks simultaneously. This leads to a more balanced view on the relative performance of different models.

We compare three variants of encoder-decoder models — including attention-based models (Bahdanau et al., 2014; Luong et al., 2015) and the model proposed by Faruqui et al. (2016) — to three very

This work is licensed under a Creative Commons Attribution 4.0 International Licence.
Licence details: http://creativecommons.org/licenses/by/4.0/

[1]We call our tasks, described below, monotone because relationships between input and output sequence characters typically obey monotonicity. That is, unlike in machine translation, there are no 'crossing edges' in corresponding alignments.

well-established baselines for monotone Seq2Seq, namely Sequitur (Bisani and Ney, 2008), DirecTL+ (Jiampojamarn et al., 2010), and Phonetisaurus (Novak et al., 2012). We also offer our own contribution[2], which may be considered a variation of the principles underlying DirecTL+. For that purpose, we have adapted higher-order pruned conditional random fields (PCRFs) (Müller et al., 2013; Lafferty et al., 2001) to handle generic monotone Seq2Seq tasks.

We find that traditional models appear to still be on par with or better than encoder-decoder models in most cases, depending on factors such as training data size and the complexity of the task at hand. We show that neural models unfold their strengths as soon as more complex phenomena need to be learned. This becomes clearly visible in the comparison between lemmatization and the other tasks we have investigated. Lemmatization is the only task at hand in which neural models outperform all established systems — as it is the only one which systematically exhibits long-range dependencies, particularly through Finnish *vowel harmony* (see Section 5). We are thus able to contrast the different challenges imposed by different tasks and show how these differences have significant impact on the performance of encoder-decoder models in comparison to established Seq2Seq models.

To our best knowledge, no systematic comparison with regard to the suitability of these encoder-decoder neural models for a wider and more generic selection of tasks has been conducted.

2 Task Description

Throughout, we denote individual tokens in a sequence by ordinary letters x, and a sequence of symbols by $\vec{x}$. Hence a string of length s is denoted as $\vec{x} = x_1 \dots x_s$. Real-valued vectors are denoted by bold-faced letters, $\mathbf{x}$.

Spelling correction is the problem of converting an 'erroneous' input sequence $\vec{x}$ into a corrected version $\vec{y}$. In terms of errors committed by humans (typos), spelling correction often deals with errors that are due to keyboard adjacency of characters and grapho-phonemic mismatch (e.g. $emergancy \rightarrow emergency$, $wuld \rightarrow would$).

OCR post-correction can be seen as a special case of spelling correction. OCR (optical character recognition) is the process of digitizing printed texts automatically, often applied to make text data from the pre-electronic age digitally available (Springmann et al., 2014). Depending on various factors including paper and scan quality, typeface, and OCR engine, OCR error rate can be extraordinarily high (Reynaert, 2014). OCR post-correction is of particular practical importance in the field of *digital humanities*. Here, paper quality, which is often bleached and tainted, and "unusual" typefaces typically cause major problems. Unlike in human spelling correction, OCR errors often arise due to *visual* similarity of character sub-sequences such as $rn \rightarrow m$ or $li \rightarrow h$.

Previous works in OCR post-correction apply noisy-channel models (Brill and Moore, 2000) and various extensions (Toutanova and Moore, 2002; Cucerzan and Brill, 2004; Gubanov et al., 2014), generic string-to-string substitution models (Xu et al., 2014), discriminative models (Okazaki et al., 2008; Farra et al., 2014), and user-interactive approaches (Reffle and Ringlstetter, 2013). Neural network designs including auto-encoders (Raaijmakers, 2013) and recurrent neural networks (Chrupała, 2014) were also investigated in previous works.

G2P conversion is the problem of converting orthographic representations into sound representations. It is the prime example of a monotone Seq2Seq task, which — as a fundamental building block for speech recognition, speech synthesis, and related tasks — has been researched for decades. It differs from the previous two tasks in that input and output strings are defined over different alphabets.

Lemmatization is the task of deriving the lemma from an inflected word form such as *atmest→atmen*. The problem is relatively simple for morphologically poor languages like English, but much harder for languages like Finnish. The task can be seen as the inverse to inflection generation (Durrett and DeNero, 2013; Ahlberg et al., 2014; Nicolai et al., 2015; Faruqui et al., 2016), where an inflected form is generated from a lemma plus an inflection tag.

[2]Our implementation of PCRF-Seq2Seq is available at:
`https://github.com/UKPLab/coling2016-pcrf-seq2seq`

3 Data

Here we detail the data sets used in our experiments; examples are provided in Table 1. These datasets reflect the different Seq2Seq tasks we aim to investigate.

The **Text+Berg** corpus (Bubenhofer et al., 2015) contains historic proceedings of the *Schweizer Alpenclub* ("Swiss Alpine Club") from the years 1864–1899 in Swiss German and French. The data has been digitized and OCR errors have been corrected manually. The corpus contains 19,024 pages, 17,186 of which are in Swiss German, and 1,838 are in French. We have extracted 88,302 unique misrecognized words along with their manually corrected counterparts.

For our experiments, we have used randomly selected 72K entries for training and test our models on another 9K entries. Furthermore, we report results for each model trained on a reduced training set (10K entries).

Twitter Typo Corpus[3]: We use a corpus of 39,172 spelling mistakes extracted from English Tweets with their respective corrections. The manually corrected mistakes come with a context word on both sides. Again, we have split the data randomly, using a training set of 31K entries and 4K for testing. We also report results on the same test set when using a reduced training set with 10K entries.

The **Combilex** data set (Richmond et al., 2009) provides mappings from English graphemes to phonetic representations (Table 1). We use different subsets for training, with 2K, 5K, 10K, and 20K entries respectively. Furthermore, we employ a test set with 26,609 entries.

P'reunde	→	**F**reunde	(misrecognition)	to_**y**ork_from	→	to_**w**ork_from
Thal wand	→	Thalwand	(segmentation)	before_**tt**_was	→	before_**i**t_was
S**lutl**erfi**m**	→	S**tud**erfi**rn**	(multiple errors)	with_my_daug**th**er	→	with_my_daug**ht**er
kinaatte	→	kinata		Waterloo	→	wOtBr5u
kinaavat	→	kinata		barnacles	→	bArn@k@5z

Table 1: Training data examples from the four corpora used: OCR detection errors for Text+Berg (top left), Twitter Typo Corpus (top right), Wiktionary Morphology Dataset (Finnish) (bottom left), and Combilex G2P mappings (bottom right).

For lemmatization, we use the **Wiktionary Morphology Dataset** (Durrett and DeNero, 2013). The data set contains inflected forms for different languages and parts of speech, corresponding lemmas, and detailed inflection information, including mood, case, and tense. We conduct experiments on the German and Finnish verb datasets and further reduce the size of the latter by considering present tense indicative verb forms in active voice only. Note that our results are not comparable to the ones presented by Durrett and DeNero (2013), Ahlberg et al. (2014), Nicolai et al. (2015), and Faruqui et al. (2016) because we focus on lemmatization, not inflection generation, as mentioned. We do so because this produces less overhead — e.g., Faruqui et al. (2016) train 27 different systems for German verbs, one for each inflection type — and there is a priori not much difference in whether we transform an inflected form to a lemma or vice versa. Hence, the relative ordering of the systems we survey should not be affected by this change of direction in the morphological analysis. In total, we have used training sets of size 43,929 entries for German verbs and 41,094 entries for Finnish verbs, and dev set and test set sizes of 5,400 (German) and 1,200 (Finnish) entries each.

4 Model Description

In this section, we briefly describe encoder-decoder neural models, pruned CRFs, and our three baselines.

4.1 Encoder-Decoder Neural Models

We compare three variants of encoder-decoder models: the 'classic' variant and two modifications:

- `enc-dec`: Encoder-decoder models using recurrent neural networks (RNNs) for Seq2Seq tasks were introduced by Cho et al. (2014) and Sutskever et al. (2014). The encoder reads an input $\vec{x}$ and

[3]Twitter typo corpus: `http://luululu.com/tweet/`

generates a vector representation $\mathbf{e}$ from it. The decoder predicts the output $\vec{y}$ one time step t at a time, based on $\mathbf{e}$. The probability for each output symbol y_t hence depends on $\mathbf{e}$ and all previously generated output symbols: $p(\vec{y}|\mathbf{e}) = \prod_{t=1}^{T'} p(y_t|\mathbf{e}, y_1 \cdots y_{t-1})$ where T' is the length of the output sequence. In NLP, most implementations of encoder-decoder models employ LSTM (long short-term memory) layers as hidden units, which extend generic RNN hidden layers with a memory cell that is able to "memorize" and "forget" features. This addresses the 'vanishing gradients' problem and allows to catch long-range dependencies.

- `attn-enc-dec`: We explore the attention-based encoder-decoder model proposed by Bahdanau et al. (2014) (Figure 1). It extends the encoder-decoder model by learning to align and translate jointly. The essential idea is that the current output unit y_t does not depend on all input units in the same way, as captured by a 'global' vector $\mathbf{e}$ encoding the input. Instead, y_t may be conditioned upon local context in the input (to which it pays *attention*).

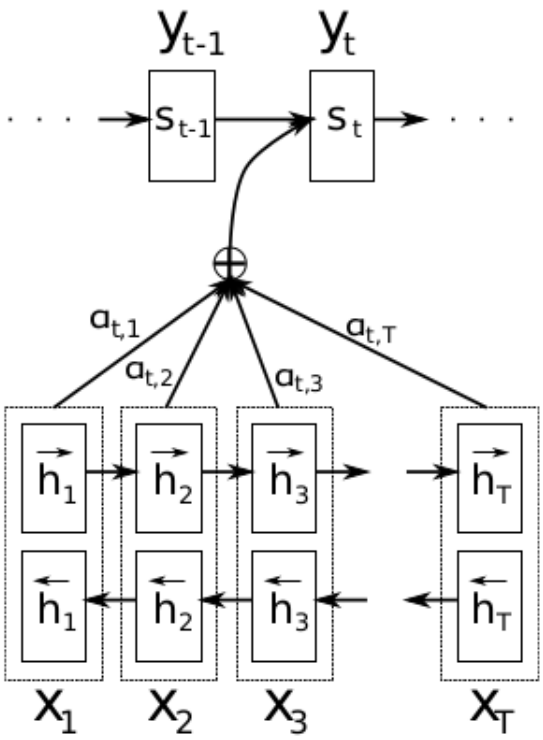

Figure 1: In the encoder-decoder model, the encoder (bottom) generates a representation of the input sequence $\vec{x}$ from which the decoder (top) generates the output sequence $\vec{y}$. The attention-based mechanism (shown here) enables the decoder to "peek" into the input at every decoding step through multiple input representations a_t. Illustration from Bahdanau et al. (2014).

- `morph-trans`: Faruqui et al. (2016) present a new encoder-decoder model designed for morphological inflection, proposing to feed the input sequence directly into the decoder. This approach is motivated by the observation that input and output are usually very similar in problems such as morphological inflection. Similar ideas have been proposed in Gu et al. (2016) in their so-called "CopyNet" encoder-decoder model (which they apply to text summarization) that allows for portions of the input sequence to be simply copied to the output sequence, without modifications. A priori, this observation seems to apply to our tasks too: at least in spelling correction, the output usually differs only marginally from the input.

For the tested neural models, we follow the same overall approach as Faruqui et al. (2016): we perform decoding and evaluation of the test data using an **ensemble** of $k = 5$ independently trained models in order to deal with the non-convex nature of the optimization problem of neural networks and the risk of

running into a local optimum (Collobert et al., 2011). The total probability p_{ens} for generating an output token y_t is estimated from the individual model output probabilities: $p_{\text{ens}}(y_t|\cdot) = \frac{1}{Z}\prod_{i=1}^{k} p_i(y_t|\cdot)^{\frac{1}{k}}$ with a normalization factor Z.

4.2 Pruned Conditional Random Fields

Conditional random fields (CRFs) were introduced by Lafferty et al. (2001) and have been a major workhorse for many sequence labeling tasks such as part-of-speech tagging and named entity recognition during the 2000s. Unfortunately, training and decoding time depend polynomially on the tag set size and exponentially on the *order* of the CRF. Here, order refers to the dependencies on the label side. This makes higher-order CRFs impractical for large training data sizes, which is the reason why virtually only first-order (linear chain) CRFs were used until recently.

Müller et al. (2013) introduced pruned CRFs (PCRFs) that approximate the CRF objective function using coarse-to-fine decoding (Charniak and Johnson, 2005). PCRFs require much shorter runtime and are thus able to make use of higher orders. Higher orders, in turn, have been shown to be highly beneficial for coarse and fine-grained part-of-speech tagging, outperforming first-order models.

For our tasks, we have adapted the implementation from Müller et al. (2013) — originally designed for sequence labeling — to general monotone Seq2Seq tasks. Sequence labeling assumes that an input sequence of length N is mapped to an output sequence of identical length N, while in Seq2Seq tasks, input string lengths may be shorter, longer, or equal to output string lengths.

We address this by first *aligning* input and output sequences as exemplified in

S	l	u	t	l	e	r	f	i	m
S	t	u	d	∅	e	r	f	i	rn

This alignment matches up character subsequences from both strings. It may include 1-to-zero matches (e.g. $l \rightarrow \emptyset$) and 1-to-many matches (e.g. $m \rightarrow rn$). We disallow many-to-1 or many-to-many matches, as they cause a problem during decoding: at test time, it is unclear how to segment a new input string into parts with size ≥ 1. A naïve 'pipeline' approach (first segment, then translate the segmented string) leads to error propagation. More sophisticated 'joint' approaches (Jiampojamarn et al., 2010) are considerably more computationally expensive.

Once the data is aligned as above, input and (modified) output sequences are of equal lengths and we can directly apply higher-order PCRFs. Below, we show that orders up to 5 (and possibly beyond) are beneficial for the Seq2Seq tasks we consider.[4] We refer to this model as **PCRF-Seq2Seq** in the remainder.

Features Conditional random fields are feature-based, so we need to decide which features we use. In view of the end-to-end nature of neural techniques, requiring little linguistic knowledge, we also minimize feature-engineering effort for the traditional approaches and thus only include very simple features. For each position p to tag, we include all consecutive character m-grams (m ranges from 1 to a maximum order of N) within a *window* of size w around p; i.e., in total our window covers $2w + 1$ positions. In our experiments below, we report results for windows of size $w = 4$ and $w = 6$. For simplicity, we set $N = w$ in each case.

4.3 Further Baseline Systems

Considering the similarity of G2P conversion, spelling correction, and lemmatization with regard to their innate monotonicity (Eger et al., 2016; Nicolai et al., 2015; Eger, 2015), we explore for all our datasets three further approaches that were originally designed for G2P conversion.

Sequitur (Bisani and Ney, 2008) is a 'joint' model for Seq2Seq in the sense of the classic distinction between joint and discriminative models. Its core architecture is a model over 'joint n-grams', also termed 'graphones' in the original publication (that is, pairs of substrings of the $\vec{x}$ and $\vec{y}$ sequence).

DirecTL+ (Jiampojamarn et al., 2010) is a discriminative model for monotone Seq2Seq that integrates joint n-gram features. It jointly learns input segmentation, output prediction, and sequence modeling.

[4] In our experiments, higher order CRFs (> 3) substantially outperformed first-order models. Typical performance differences were from about 4 to 7% between first-order and fifth-order models. For brevity, we omit results for orders ≤ 3.

	Text+Berg (72K)	**Text+Berg (10K)**	**Twitter (31K)**	**Twitter (10K)**
attn-enc-dec (1 layer, size 100)	66.80%	61.71%	66.25%	60.99%
attn-enc-dec (1 layer, size 200)	68.29%	**63.00%**	67.81%	59.36%
attn-enc-dec (2 layers, size 100)	68.30%	62.27%	69.29%	63.31%
attn-enc-dec (2 layers, size 200)	69.74%	62.87%	69.70%	62.01%
enc-dec (1 layer, size 100)	50.96%	39.99%	60.91%	52.90%
enc-dec (1 layer, size 200)	53.65%	41.52%	63.39%	55.76%
enc-dec (2 layers, size 100)	56.94%	42.53%	65.94%	56.50%
enc-dec (2 layers, size 200)	59.01%	46.01%	63.70%	53.74%
morph-trans (1 layer, size 100)	55.96%	49.37%	41.46%	39.19%
morph-trans (1 layer, size 200)	54.22%	49.63%	36.28%	30.74%
morph-trans (2 layers, size 100)	56.11%	47.35%	44.42%	42.02%
morph-trans (2 layers, size 200)	49.27%	45.55%	30.33%	29.61%
PCRF-Seq2Seq (order 4, $w = 6$)	**74.67%**	62.24%	73.52%	59.97%
PCRF-Seq2Seq (order 5, $w = 6$)	74.22%	62.47%	74.03%	60.19%
PCRF-Seq2Seq (order 4, $w = 4$)	74.55%	62.75%	**74.87%**	**63.59%**
Phonetisaurus ($n = 8$)	60.89%	51.84%	69.52%	55.76%
Sequitur	68.04%	57.30%	70.74%	58.90%

Table 2: Word accuracies (WACs) for all encoder-decoder models, PCRF-Seq2Seq, and baselines for the **OCR post-correction** task and for the **spelling correction** task. Best configurations for each model are underlined, overall best results are bold-faced.

Since it is based on ordinary CRFs, it is virtually impossible to use this system with higher orders for all practically relevant datasets due to very long training times. Moreover, the system is generally very slow because it jointly learns to segment and translate, as mentioned. For this reason, we have only tested it on the Combilex dataset (Table 3), run with comparable parametrizations (context size, etc.) as PCRF-Seq2Seq.

Phonetisaurus (Novak et al., 2012) implements a weighted finite state transducer (WFST) to align input and output tokens. The EM-driven algorithm is capable of learning multiple-to-multiple alignments where we restrict both sides to a maximum of 2. The alignments learned from the training data are subsequently used to train a character-based n-gram language model. For brevity, we only report results for models with $n = 8$, which outperformed lower-order models in our experiments.

5 Results and Analysis

5.1 Model Performances

We report the results of all our experiments in terms of *word accuracy* (WAC), i.e., the fraction of completely correctly predicted output sequences. Table 2 lists WACs for all our systems on the OCR post-correction task (*Text+Berg*, full and reduced training set) and on the spelling correction task (Twitter, full and reduced training set). Table 3 reports WAC of all tested models on the Combilex dataset with models trained on training sets of different sizes. Table 4 reports WAC for the lemmatization task on the morphology dataset.

For the encoder-decoder models, we report the results with one and two layers of sizes 100 and 200 each. We have additionally conducted sample experiments with larger networks which have shown that neither increasing the number of layers nor the size of the layers leads to further improvements. For the PCRF-Seq2Seq models, we report results for windows of sizes $w = 4$ and $w = 6$. We note that, a priori, more training data tends to favor larger context size w, whereas a large w may lead to overfitting when training data is small. The same holds for model order.

	Combilex (20K)	**Combilex (10K)**	**Combilex (5K)**	**Combilex (2K)**
attn-enc-dec (1 layer, size 100)	57.39%	54.40%	41.19%	35.68%
attn-enc-dec (1 layer, size 200)	61.86%	57.92%	46.72%	38.31%
attn-enc-dec (2 layers, size 100)	66.74%	60.52%	55.89%	44.13%
attn-enc-dec (2 layers, size 200)	67.36%	59.62%	55.07%	44.26%
enc-dec (1 layer, size 100)	54.03%	48.25%	36.62%	18.17%
enc-dec (1 layer, size 200)	55.27%	49.81%	36.19%	18.41%
enc-dec (2 layers, size 100)	57.77%	51.91%	38.97%	16.68%
enc-dec (2 layers, size 200)	56.95%	50.69%	39.01%	17.83%
morph-trans (1 layer, size 100)	48.82%	43.30%	33.63%	18.97%
morph-trans (1 layer, size 200)	49.72%	43.15%	32.42%	18.76%
morph-trans (2 layers, size 100)	49.58%	42.05%	28.69%	15.13%
morph-trans (2 layers, size 200)	44.36%	35.14%	23.08%	12.74%
PCRF-Seq2Seq (order 4, $w = 6$)	72.14%	64.39%	55.66%	42.82%
PCRF-Seq2Seq (order 5, $w = 6$)	72.23%	64.32%	55.58%	42.62%
PCRF-Seq2Seq (order 4, $w = 4$)	71.74%	64.71%	**56.89%**	**44.74%**
DirecTL+	72.23%	**65.09%**	55.75%	42.95%
Phonetisaurus ($n = 8$)	**72.29%**	64.14%	55.28%	42.21%
Sequitur	70.57%	62.57%	54.03%	41.94%

Table 3: WACs for all encoder-decoder models, PCRF-Seq2Seq, and baselines for the **G2P** task. Best configurations for each model are underlined, overall best results are bold-faced.

While more training data obviously increases WAC for every model, the specific impact varies. In general, (attention-based) encoder-decoder models deal relatively well with limited test data in our experiments, achieving WACs comparable to PCRF-Seq2Seq. In contrast, they appear to benefit less from increasing data sizes than CRFs do. On the Twitter dataset, for instance, the best-performing encoder-decoder model increases WAC by 7.7 percentage points when tripling the training data size. At the same time, the 5th-order PCRF-Seq2Seq WAC increases by 13.8. When a large amount of training data is available, CRFs therefore consistently outperform neural models, and so do the specialized baseline systems on the G2P conversion tasks (Table 3).

Summarizing, we find that PCRF-Seq2Seq performs best among the tested systems for the two spelling correction tasks when large training data is available. The best performance of PCRF-Seq2Seq is roughly 6-7 percentage points better than the best performance of an encoder-decoder model for both Twitter 31K and Text+Berg 72K. For small training set sizes, PCRF-Seq2Seq and the encoder-decoder models are on a similar level. For the G2P task, an analogous pattern emerges. Moreover, here, all classical systems appear to perform similarly, with DirecTL+ and PCRF-Seq2Seq marginally outperforming the others. For lemmatization, the overall picture looks different. For Finnish verbs, we observe the only case in which attention-based encoder-decoder systems clearly outperform all other approaches. For German, neural models also achieve the best results, albeit only marginally above PCRF-Seq2Seq.

Previous works that employed encoder-decoder models successfully focused on tasks like machine translation and grammar correction in which more challenging linguistic phenomena such as long-range dependencies and 'crossing edges' (re-ordering) occur frequently. In our experiments, too, neural models only outperform traditional ones when long-range dependencies become relevant, namely in lemmatization. In all other tasks at hand, in contrast, neural models perform worse or equal.

The afore-mentioned, more complex linguistic phenomena intuitively require a more global view on long input sequences which is hard to impossible to model for approaches that cannot look beyond a statically defined context. Spelling mistakes, OCR errors, and G2P, however, largely depend on a very local context. For instance, OCR systems typically do not consider more than a small context when

	German Verbs (44K)	Finnish Verbs (41K)
attn-enc-dec (1 layer, size 100)	93.67%	**98.00%**
attn-enc-dec (1 layer, size 200)	93.17%	96.92%
attn-enc-dec (2 layers, size 200)	92.39%	97.08%
attn-enc-dec (2 layers, size 200)	**94.83%**	96.42%
enc-dec (1 layer, size 100)	77.31%	94.83%
enc-dec (1 layer, size 200)	82.00%	94.67%
enc-dec (2 layers, size 200)	79.50%	95.67%
enc-dec (2 layers, size 200)	76.30%	95.58%
morph-trans (1 layer, size 100)	91.89%	96.17%
morph-trans (1 layer, size 200)	93.24%	96.75%
morph-trans (2 layers, size 200)	93.02%	97.08%
morph-trans (2 layers, size 200)	93.63%	96.75%
PCRF-Seq2Seq (order 4, $w = 6$)	94.22%	94.08%
PCRF-Seq2Seq (order 5, $w = 6$)	93.77%	94.00%
PCRF-Seq2Seq (order 4, $w = 4$)	93.44%	93.33%
Phonetisaurus ($n = 8$)	86.62%	93.42%
Sequitur	85.63%	92.92%

Table 4: WACs for all encoder-decoder models, PCRF-Seq2Seq, and baselines for the **lemmatization task**. Best configurations for each model are underlined, overall best results are bold-faced.

estimating the probability of a character. Regarding the G2P task, phonetics is generally independent of characters that occur more than two or three positions before of after, at least in most cases and in English. The same is presumably true for human typos, where a mistaken key stroke may be the result of a previous key's position, but does not correlate to any key that was hit several time steps before. Hence, neural networks are unable to benefit from their often advantageous capability of modeling long-range dependencies here.

Especially the Finnish lemmatization experiments confirm that the capability of dealing with long-range dependencies plays an important role. Finnish *vowel harmony* makes a vowel control other vowels in the word, potentially across multiple syllables; see Faruqui et al. (2016) for more detailed explanation.

As a side note, our results are in line with the common notion that the specific impact of a neural network's size (number and sizes of layers) is almost unpredictable. Our results can only confirm the general rule-of-thumb that larger networks are better for larger training sets, while models with fewer parameters outperform larger ones when training data is smaller.

5.2 Training Time

Another potentially limiting factor for the applicability of a model in real-world scenarios, especially for large datasets, is training time. Under all circumstances, weighted finite state transducers (Phonetisaurus) are trained by magnitudes faster than all other approaches. Training times range between as little as 6 seconds for the smallest training set and 247 seconds for the full *Text+Berg* training set (72K entries).

In comparison, training times for the encoder-decoder models range from 2 to 80 hours (without using GPUs) for 30 epochs, depending on the sizes of the networks and the training data. Furthermore, there is no noticeable difference between either of the three encoder-decoder variations. Training time increases approximately linearly with the number of layers, the size of the layers, and the training data size. All these factors add up, meaning that doubling both the number of layers and the size of the layers approximately quadruples training time.

Contrasting DirecTL+ with PCRF-Seq2Seq, both of which rest on similar principles and also perform similarly in our experiments on the G2P task, we find that training PCRF-Seq2Seq was a factor of 30 or 50 times faster than DirecTL+ on Combilex (2K) and Combilex (5K), respectively. In general, training

for PCRF-Seq2Seq across our datasets was in the order of minutes to (few) hours.

5.3 Error Analysis

We divided three of our test sets (*Text+Berg*, Twitter, and Combilex) by input string lengths and evaluated PCRF-Seq2Seq and encoder-decoder neural models on these subsets of the test data. As illustrated in Figures 2 and 3, we observe a consistent tendency: PCRF-Seq2Seq performs relatively robustly over input strings of different lengths, while the performance of the encoder-decoder models plummets more drastically with sequences becoming longer, in particular those without attention-mechanism.

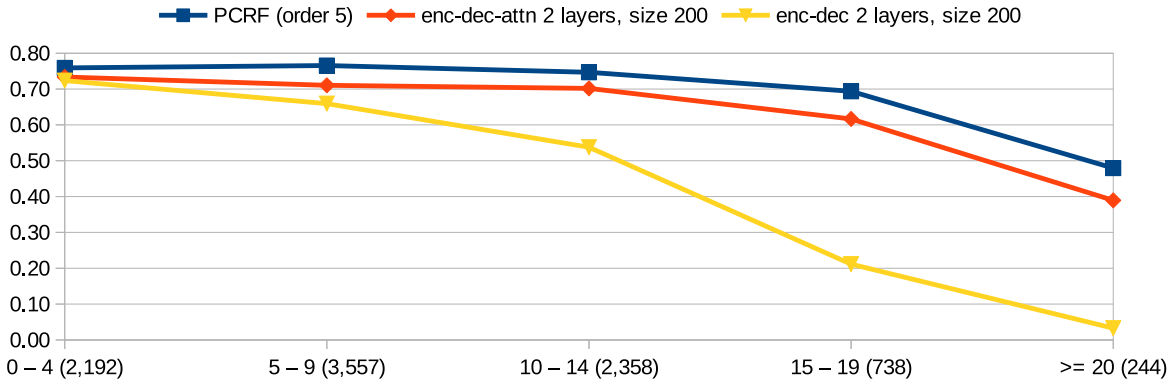

Figure 2: WAC of PCRF-Seq2Seq and encoder-decoder neural models with and without attention-based mechanisms as a function of input string length (number of training samples) on *Text+Berg* **OCR post-correction**.

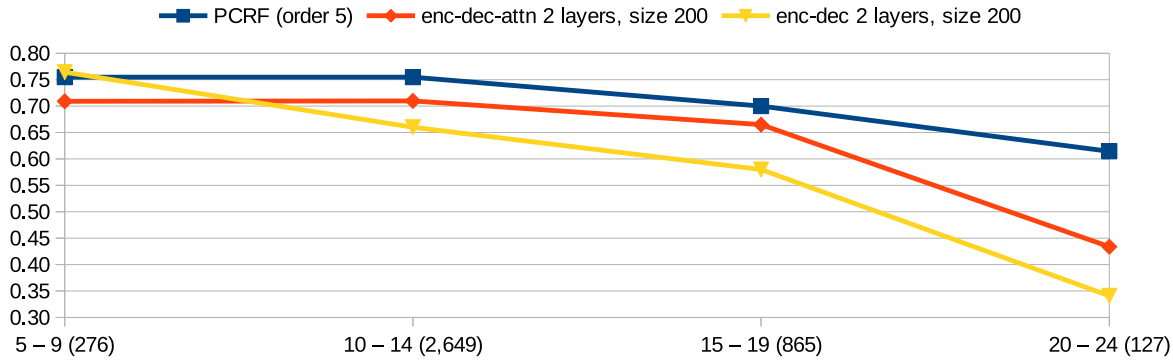

Figure 3: WAC of PCRF-Seq2Seq and encoder-decoder neural models with and without attention-based mechanisms as a function of input string length (number of training samples) on the Twitter **spelling correction** task.

For shorter sequences, we observe that standard encoder-decoder models even slightly outperform their attention-based counterparts as well as PCRF-Seq2Seq on both the Twitter spelling correction task (Figure 3) and on G2P conversion, in contrast to their rather low performance on the full datasets. On the *Text+Berg* data, all systems achieve approximately equal WAC for short sequences (Figure 2).

For longer sequences, the performance of the encoder-decoder models drops dramatically on all data sets. This effect is also visible, albeit less strong, for the attention-based variant. This can be seen particularly well on the OCR post-correction task (Figure 2), where the test set contains numerous long sequences: the accuracy rate for the standard encoder-decoder model drops from 73.32% on very short sequences to below 10% for very long ones ($\geq$ 20), whereas the attention-based model drops less drastically to 38.93% (from 76.92%). At the same time, PCRF-Seq2Seq behaves more stably, particularly on the Twitter data (Figure 3). For the Combilex data, the picture looks very similar — we omit these results for brevity.

6 Conclusions

The generality of neural networks makes them appealing for a wide range of possible tasks. In the scope of this work, we have applied encoder-decoder neural models to monotone Seq2Seq tasks. We have shown that they can perform comparably to more specialized models in some cases, but cannot (yet) consistently outperform established approaches, and are sometimes still substantially below them. Furthermore, the advantage of having rendered feature engineering and hyper-parameter optimization in the traditional sense unnecessary is notoriously substituted by the search for optimal neural network topologies.

At first sight, our analyses based on string lengths are in line with those reported by Bahdanau et al. (2014). They state that — for the field of machine translation — the attention mechanism leads to improvements over the standard encoder-decoder model on longer sentences. We also observe this positive impact for our tasks, where the attention-based mechanism alleviates the drastic performance drop of the standard encoder-decoder models on long sequences to some extent. At the same time, we see that very performance drop persisting — CRFs still outperform encoder-decoder models on long sequences, even when employing attention-mechanisms. As described in Section 5.3, neural models are only able to successfully compete when more complex phenomena occur, on which traditional models fail. Nevertheless, previous works such as Vukotic et al. (2015) also indicate that even in more complex sequence labeling tasks such as spoken language understanding, neural networks are not guaranteed to outperform CRFs.

The task-specific extensions to the encoder-decoder proposed by Faruqui et al. (2016) have been shown to produce mostly bad results in our settings. This is particularly surprising for the OCR data, for which input and output sequences are usually very similar, so that we had expected that re-feeding the input to the decoder should be equally beneficial in that domain. As discussed, one explanation might be that OCR, or spelling correction generally, putatively exhibits few long-range dependencies. This might explain why the `morph-trans` approach works quite well and competitive in morphological analysis tasks, as re-confirmed in our experiments. Thus, long-range dependencies might actually be a more crucial aspect for the performance of the model presented by Faruqui et al. (2016) than the similarity between input and output sequence.

We conclude that neural networks are far from completely replacing established methods at this point, as the latter can be both faster and more accurate, depending on the properties of the task at hand. A systematic analysis of the complexities and challenges a particular task imposes, remains unavoidable. At the same time, one can argue that encoder-decoder neural models are a relatively recent development and might continue to improve much over the next years. Being very generic and largely task-agnostic, they are already able to outperform traditional and specialized approaches under certain circumstances.

Acknowledgements

This work has been supported by the Volkswagen Foundation as part of the Lichtenberg-Professorship Program under grant № I/82806, and by the German Institute for Educational Research (DIPF), as part of the graduate program ”Knowledge Discovery in Scientific Literature“ (KDSL).

References

Malin Ahlberg, Markus Forsberg, and Mans Hulden. 2014. Semi-supervised learning of morphological paradigms and lexicons. In *Proceedings of the 14th Conference of the European Chapter of the Association for Computational Linguistics, Gothenburg, Sweden 26–30 April 2014*, pages 569–578.

Dzmitry Bahdanau, Kyunghyun Cho, and Yoshua Bengio. 2014. Neural Machine Translation by Jointly Learning to Align and Translate. *arXiv:1409.0473 [cs, stat]*, September.

Maximilian Bisani and Hermann Ney. 2008. Joint-Sequence Models for Grapheme-to-Phoneme Conversion. *Speech Communication*, 50(5):434–451, May.

Eric Brill and Robert C. Moore. 2000. An Improved Error Model for Noisy Channel Spelling Correction. In *Proceedings of ACL ’00*, pages 286–293, Stroudsburg, PA, USA. Association for Computational Linguistics.

Noah Bubenhofer, Martin Volk, Fabienne Leuenberger, and Daniel Wüest, editors. 2015. *Text+Berg-Korpus (Release 151_v01)*. Institut für Computerlinguistik, Universität Zürich. Digitale Edition des Jahrbuch des SAC 1864–1923, Echo des Alpes 1872–1924, Die Alpen, Les Alpes, Le Alpi 1925–2014, The Alpine Journal 1969–2008. Published: XML-Format.

Eugene Charniak and Mark Johnson. 2005. Coarse-to-Fine n-Best Parsing and MaxEnt Discriminative Reranking. In *Proceedings of ACL '05*, pages 173–180, Ann Arbor, MI, USA. Association for Computational Linguistics.

Kyunghyun Cho, Bart van Merrienboer, Caglar Gulcehre, Dzmitry Bahdanau, Fethi Bougares, Holger Schwenk, and Yoshua Bengio. 2014. Learning Phrase Representations using RNN Encoder-Decoder for Statistical Machine Translation. *arXiv:1406.1078 [cs, stat]*, June.

Grzegorz Chrupała. 2014. Normalizing Tweets with Edit Scripts and Recurrent Neural Embeddings. In *Proceedings of ACL '14*, pages 680–686, Baltimore, MD, USA. Association for Computational Linguistics.

Ronan Collobert, Jason Weston, Léon Bottou, Michael Karlen, Koray Kavukcuoglu, and Pavel Kuksa. 2011. Natural Language Processing (Almost) from Scratch. *The Journal of Machine Learning Research*, 12:2493–2537, February.

Silviu Cucerzan and Eric Brill. 2004. Spelling Correction as an Iterative Process that Exploits the Collective Knowledge of Web Users. In *Proceedings of EMNLP '04*, pages 293–300, Barcelona, Spain. Association for Computational Linguistics.

Greg Durrett and John DeNero. 2013. Supervised learning of complete morphological paradigms. In *Proceedings of NAACL-HLT '13*, pages 1185–1195, Atlanta, GA, USA. Association for Computational Linguistics.

Steffen Eger, Tim vor der Brück, and Alexander Mehler. 2016. A Comparison of Four Character-Level String-to-String Translation Models for (OCR) Spelling Error Correction. *The Prague Bulletin of Mathematical Linguistics*, 105(1):77–99, April.

Steffen Eger. 2015. Designing and comparing g2p-type lemmatizers for a morphology-rich language. In *Systems and Frameworks for Computational Morphology - Fourth International Workshop, SFCM 2015, Stuttgart, Germany, September 17-18, 2015, Proceedings*, pages 27–40.

Noura Farra, Nadi Tomeh, Alla Rozovskaya, and Nizar Habash. 2014. Generalized Character-Level Spelling Error Correction. In *Proceedings of ACL '14*, pages 161–167, Baltimore, MD, USA. Association for Computational Linguistics.

Manaal Faruqui, Yulia Tsvetkov, Graham Neubig, and Chris Dyer. 2016. Morphological Inflection Generation Using Character Sequence to Sequence Learning. In *Proceedings of NAACL-HLT '16*, pages 634–643, San Diego, CA, USA. Association for Computational Linguistics.

Jiatao Gu, Zhengdong Lu, Hang Li, and Victor O. K. Li. 2016. Incorporating copying mechanism in sequence-to-sequence learning. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics, ACL 2016, August 7-12, 2016, Berlin, Germany, Volume 1: Long Papers*.

Sergey Gubanov, Irina Galinskaya, and Alexey Baytin. 2014. Improved Iterative Correction for Distant Spelling Errors. In *Proceedings of ACL '14*, pages 168–173, Baltimore, MD, USA. Association for Computational Linguistics.

Sittichai Jiampojamarn, Colin Cherry, and Grzegorz Kondrak. 2010. Integrating Joint n-gram Features into a Discriminative Training Framework. In *Proceedings of NAACL-HLT '10*, pages 697–700, Los Angeles, CA, USA. Association for Computational Linguistics.

John Kominek and Alan W Black. 2004. The CMU Arctic speech databases. In *Fifth ISCA Workshop on Speech Synthesis*, pages 223–224, Pittsburgh, PA, USA.

John Lafferty, Andrew McCallum, and Fernando Pereira. 2001. Conditional Random Fields: Probabilistic Models for Segmenting and Labeling Sequence Data. In *Proceedings of ICML '01*, pages 282–289, San Francisco, CA, USA. Morgan Kaufmann Publishers Inc.

Mark Lewellen. 1998. Neural Network Recognition of Spelling Errors. In *Proceedings of ACL-COLING '98*, pages 1490–1492, Montréal, Québec, Canada. Association for Computational Linguistics.

Thang Luong, Hieu Pham, and Christopher D. Manning. 2015. Effective Approaches to Attention-based Neural Machine Translation. In *Proceedings of EMNLP '15*, pages 1412–1421, Lisbon, Portugal. Association for Computational Linguistics.

Thomas Müller, Helmut Schmid, and Hinrich Schütze. 2013. Efficient Higher-Order CRFs for Morphological Tagging. In *Proceedings of EMNLP '13*, pages 322–332, Seattle, WA, USA. Association for Computational Linguistics.

Garrett Nicolai, Colin Cherry, and Grzegorz Kondrak. 2015. Inflection generation as discriminative string transduction. In *Proceedings of NAACL-HLT '15*, pages 922–931, Denver, CO, USA. Association for Computational Linguistics.

Josef R. Novak, Nobuaki Minematsu, and Keikichi Hirose. 2012. WFST-Based Grapheme-to-Phoneme Conversion: Open Source tools for Alignment, Model-Building and Decoding. In *Proceedings of the 10th International Workshop on Finite State Methods and Natural Language Processing*, pages 45–49, Donostia, Spain. Association for Computational Linguistics.

Naoaki Okazaki, Yoshimasa Tsuruoka, Sophia Ananiadou, and Jun'ichi Tsujii. 2008. A Discriminative Candidate Generator for String Transformations. In *Proceedings of EMNLP '08*, pages 447–456, Honolulu, HI, USA. Association for Computational Linguistics.

Stephan Raaijmakers. 2013. A Deep Graphical Model for Spelling Correction. In *Proceedings of the 25th Benelux Conference on Artificial Intelligence*, pages 160–167, Delft, Netherlands. Delft University of Technology.

Kanishka Rao, Fuchun Peng, Hasim Sak, and Françoise Beaufays. 2015. Grapheme-to-Phoneme Conversion Using Long Short-Term Memory Recurrent Neural Networks. In *Proceedings of ICASSP '15*, pages 4225–4229, South Brisbane, QLD, Australia. Institute of Electrical and Electronics Engineers.

Ulrich Reffle and Christoph Ringlstetter. 2013. Unsupervised Profiling of OCRed Historical Documents. *Pattern Recognition*, 46(5):1346–1357, May.

Martin Reynaert. 2014. On OCR Ground Truths and OCR Post-correction Gold Standards, Tools and Formats. In *Proceedings of DATeCH '14*, pages 159–166, New York, NY, USA. ACM.

Korin Richmond, Robert A.J. Clark, and Susan Fitt. 2009. Robust LTS rules with the Combilex speech technology lexicon. In *Proceedings of INTERSPEECH '09*, pages 1295–1298, Brighton, UK. ISCA.

Allen Schmaltz, Yoon Kim, Alexander M Rush, and Stuart M Shieber. 2016. Sentence-Level Grammatical Error Identification as Sequence-to-Sequence Correction. *arXiv preprint arXiv:1604.04677*.

Tarek Sherif and Grzegorz Kondrak. 2007. Substring-Based Transliteration. In *Proceedings of ACL '07*, pages 944–951, Prague, Czech Republic. Association for Computational Linguistics.

Uwe Springmann, Dietmar Najock, Hermann Morgenroth, Helmut Schmid, Annette Gotscharek, and Florian Fink. 2014. OCR of Historical Printings of Latin Texts: Problems, Prospects, Progress. In *Proceedings of the First International Conference on Digital Access to Textual Cultural Heritage*, pages 71–75, New York, NY, USA. ACM.

Ilya Sutskever, Oriol Vinyals, and Quoc VV Le. 2014. Sequence to Sequence Learning with Neural Networks. In *Proceedings of Neural Information Processing Systems 2014*, pages 3104–3112, Montréal, Québec, Canada. Curran Associates, Inc.

Kristina Toutanova and Robert C. Moore. 2002. Pronunciation modeling for improved spelling correction. In *Proceedings of ACL '02*, pages 144–151, Philadelphia, PA, USA. Association for Computational Linguistics.

Oriol Vinyals and Quoc Le. 2015. A Neural Conversational Model. In *Proceedings of the 31st International Conference on Machine Learning, JMLR: W&CP*, volume 37, Lille, France.

Vedran Vukotic, Christian Raymond, and Guillaume Gravier. 2015. Is it time to switch to Word Embedding and Recurrent Neural Networks for Spoken Language Understanding? In *InterSpeech-2015*, pages 130–134, Dresden, Germany.

Ziang Xie, Anand Avati, Naveen Arivazhagan, Dan Jurafsky, and Andrew Y Ng. 2016. Neural Language Correction with Character-Based Attention. *arXiv preprint arXiv:1603.09727*.

Gu Xu, Hang Li, Ming Zhang, and Ziqi Wang. 2014. A Probabilistic Approach to String Transformation. *IEEE Transactions on Knowledge and Data Engineering*, 26(5):1063–1075, May.

Kaisheng Yao and Geoffrey Zweig. 2015. Sequence-to-Sequence Neural Net Models for Grapheme-to-Phoneme Conversion. In *Proceedings of INTERSPEECH '15*, pages 3330–3334, Dresden, Germany. ISCA.

Wenpeng Yin, Sebastian Ebert, and Hinrich Schütze. 2016. Attention-Based Convolutional Neural Network for Machine Comprehension. In *Proceedings of the Workshop on Human-Computer Question Answering*, pages 15–21, San Diego, CA, USA. Association for Computational Linguistics.